Generalized spatial and temporal interpolation, and limited data reconstruction

Bruno Madore, Ph.D.
Harvard Medical School, Brigham and Women's Hospital, Department of Radiology, Boston, USA

1. INTRODUCTION

1.1 Overview: In this course, a series of fast-imaging techniques will be described. These techniques are designed to reduce the amount of data that needs to be acquired, leading to shorter scan times.

1.2 Traditional approaches toward fast imaging: Traditional ways of increasing speed have involved pulse sequence design, and hardware developments. As described in the 'physics of fast imaging' session, imaging sequences such as EPI, RARE and SSFP are fast because they were designed to spend as much of the available time as possible actually acquiring data, to minimize idle time. On the hardware front, the importance of having good-performance gradient sets can be appreciated through an analogy with auto racing. The acquisition process involves 'driving' through all the needed data-point locations, in *k*-space. The maximum gradient strength corresponds to the maximum speed of the car, while the gradient slew-rate corresponds to acceleration, braking and steering. Accordingly, strong gradients that can be switched on and off rapidly allow fast racing through *k*-space, and short scan times. While fast imaging sequences offer efficient itineraries through *k*-space, with little idle time, it is the hardware, especially gradient sets, that determines how fast one can navigate around.

1.3 Going faster, by skipping data points: Running fast-imaging pulse sequences on modern MR systems, with good-performance gradient sets, enables fast imaging. But technology and safety impose limits on achievable hardware performance, and physics imposes limits on how efficient a pulse sequence can be. Once traditional approaches have reached their limit and further imaging speed is still required, more drastic solutions have to be considered. These solutions involve further increasing imaging speed by skipping some of the data that should normally be acquired. The skipped data can be calculated, based on the data we do acquire, using assumptions/prior knowledge. By calculating instead of acquiring parts of the dataset, the corresponding '*k*-space itinerary' can be shortened, and acquisition time is reduced. These partially sampled approaches are able to carry imaging speed many-fold beyond the limits technology and physics would seem, at first sight, to impose.

1.4 Interpolation / limited data reconstruction: The fast-imaging methods considered here involve reconstructing limited datasets, i.e., datasets that were only partly sampled. The reconstruction can be described as an interpolation process, in the sense that missing data points are calculated based on a collection of known data points. But unlike traditional interpolation algorithms, often based on simple kernels such as sinc or linear functions, most methods discussed here exploit information specific to the MR imaging process, and/or to the imaged object itself. The overall value of these methods can be judged based on their ability to reduce data requirements (ratio of the size of the full and limited data sets, sometimes called "acceleration factor"), as well as on the amount of artifacts and noise they may introduce doing so.

1.5 Many different ways of skipping data: As a reference, Fig. 1a depicts a fully sampled *k*-space matrix. Although in reality *k*-space matrices typically feature about one or two hundred lines, only 12 lines are actually depicted in Fig. 1, to keep the drawings visually simple. Figures 1b through 1f depict a number of different ways to skip data, to speed-up the image acquisition process. Full black lines represent acquired data, while gray lines represent locations that are skipped. In the following sections, we will review a number of methods based on these sub-sampling strategies. Special emphasis will be placed on the

assumption(s) and/or prior knowledge used for reconstruction, on situations where such assumption(s)/knowledge may prove inaccurate, and on the acceleration achieved. While data can be skipped in many different ways (Fig. 1), the main challenge, of course, lies in recovering the missing information.

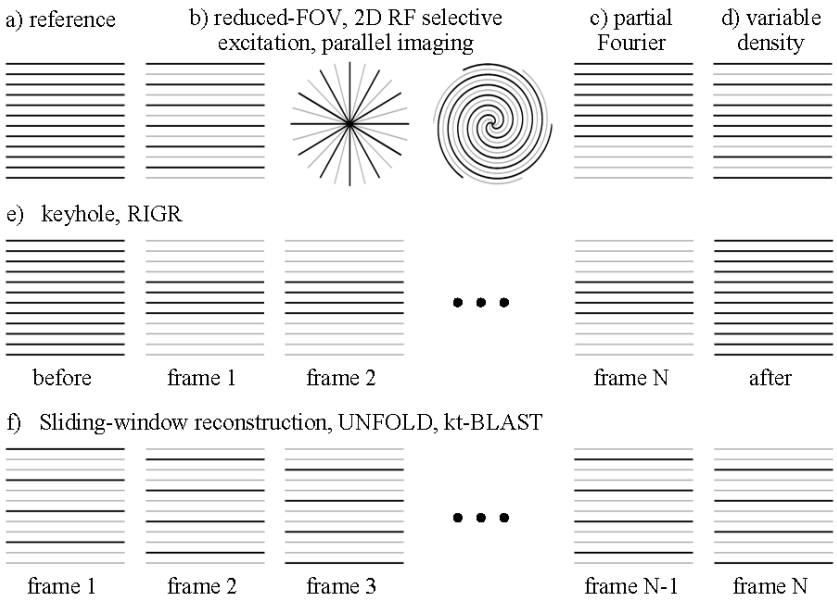

Fig. 1: A fully sampled *k*-space matrix is depicted in (a). Data can be sub-sampled in a regular fashion, throughout *k*-space (b). Most sub-sampling strategies can, typically, be extended to non-Cartesian imaging as well. Examples for radial and spiral sampling schemes are depicted here. Alternately, one can fully sample one side of *k*-space, and skip most of the other side (c). Or based on an argument that the central region is most important, one can vary the sampling density across *k*-space (d). In dynamic applications, when a time series of images is acquired, one has a new dimension available, and different sampling strategies can be chosen for different time frames (e and f).

2. DESCRIPTION OF INDIVIDUAL TECHNIQUES

2.1 Partial-Fourier (Fig. 1c): MR images are complex, in the sense that each pixel features both a magnitude and a phase value. The magnitude is generally the quantity displayed for diagnostic purposes, while the phase is often discarded. MR signal comes from magnetization in the transverse plane, and the phase of the MR signal corresponds to the orientation of the magnetization vector. In theory, RF pulses that flip the magnetization onto the "real" axis in the transverse plane should lead to entirely real images, associated with k-space matrices of Hermitian symmetry. In this case, there would be no point in acquiring both halves of k-space: one could acquire a single half, assume k-space to have Hermitian symmetry, and calculate the missing half from the acquired one. As a result, a two-fold reduction in imaging time would be achieved. Unfortunately, many imperfections (e.g. magnetic susceptibility, field inhomogeneities) can affect the phase of the signal, leading to non-Hermitian k-space data. In practical implementations of partial-Fourier, the requirement for a real object is softened, and the phase is allowed to vary slowly spatially (instead of zero everywhere). The softer assumption leads to smaller decreases in data

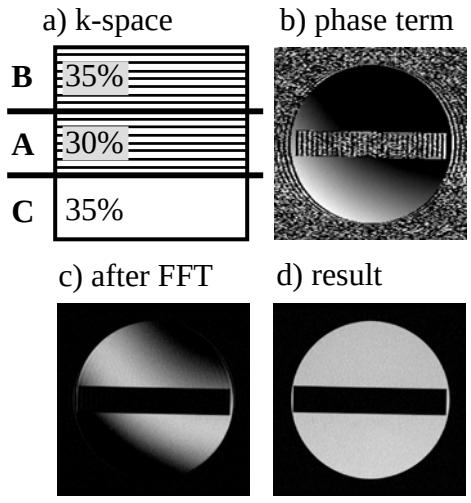

Fig. 2: Homodyne reconstruction, where 65% of a full k-space matrix is acquired. a) A fully acquired region **A,** covering 30% of the full matrix and symmetrical with respect to the k-space origin, is flanked with two equally sized regions: A fully sampled region **B** and a non-sampled region **C**. b) Data from region **A** is FFTed, and the phase of the result is displayed here. The phase-encoding direction is horizontal in the present example. c) A composite k-space matrix is created as follows. Region **C** is filled with zeros, data from region **A** is pasted in, and data from region **B**, multiplied by 2, is also pasted in. This composite matrix is FFTed, and the real component of the result is shown here, in (c). d) The final result is obtained by using the data in (b) to phase-correct the complex data from (c). After the phase correction, the desired signal is found in the real component, shown in (d), while the imaginary component is expected to contain only noise and artifacts, and is simply discarded.

requirements, i.e., more than 50% of the k-space matrix needs to be acquired. A few different implementations have been proposed (e.g., [1-5]), and the "homodyne" reconstruction [3] approach is illustrated in Fig. 2, as an example.

Because the method assumes a spatially smooth phase, artifacts may occur near locations where the spatial derivative of the magnetic field is high, e.g., near a piece of metal, or near a tissue-air transition. Partial-Fourier can be applied either in the phase- or frequency-encoding direction (but not both at once). Partial-Fourier in the phase-encoding direction generally leads to a decrease in scan time, as the number of acquired k_y lines is reduced (Fig. 2). Partial-Fourier in the frequency-encoding direction allows TE and the amount of gradient switching before the readout to be reduced, resulting in sequences more robust to intravoxel dephasing and signal loss, and thus especially useful in flow applications.

2.2 Variable-density imaging (Fig. 1d): In variable-density imaging, the outer regions of k-space are sampled more sparsely than the center [6,7,8]. While the central region is sampled densely enough to provide an FOV that is at least as large as the imaged object (to avoid aliasing artifacts), the outer regions are sampled more coarsely in order to save time. Sampling data inappropriately in the outer regions of k-space is expected to cause artifacts, but because most of the k-space energy is found near the (correctly sampled) center, these artifacts are expected to be faint. For non-Cartesian sampling functions such as projection-reconstruction (PR) and spiral imaging, aliasing artifacts tend to be non-structured, and smeared over the whole FOV. Accordingly, the faint aliasing artifacts from object edges are further "diluted" by being spread over a large spatial area, and typically appear noise-like in the resulting images [6].

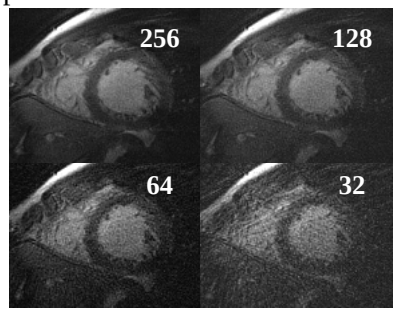

Fig. 3: Courtesy of Dr. Dana C. Peters, Beth Israel Deaconess Medical Center, Boston, USA. A PR cardiac image is reconstructed with 256, 128, 64 and 32 projections. Notice that reconstructing with only 128 or 64 projections, instead of the full 256, has only a small effect on artifact content in this case.

2.3 Parallel imaging (Fig. 1b): Parallel imaging uses prior knowledge about the sensitivity of the imaging coils to evaluate missing data. This popular technique was described as part of Dr. Larkman's tutorial, and will not be further treated here.

2.4 2D RF selective excitations (Fig. 1b): The Fourier transform operation is non-localized, meaning that any spatial location within the object contributes to all locations in the Fourier domain. In other words, even if one is interested only in a small portion of the object, all k-space locations must nevertheless be obtained. Unlike frequency- and phase-encoding processes based on Fourier functions, the slice selection applied along the z direction is a localized operation. If one is interested in a specific z location, an x-y plane can be excited for this z location only. Using more complicated RF pulses, able to selectively excite along two directions instead of only one [9,10], RF encoding can be used to assist phase encoding, and speed-up the imaging process. In its simplest form, 2D selective excitation could be used to excite only 1/*n* of the object in the phase-encoding (y) direction, allowing an *n*-fold reduction in FOV size and in the number of k-space lines required. However, the minimum length of a 2D selective RF pulse is normally longer than that of a simpler 1D pulse, leading to longer minimum values for TR and TE. Furthermore, relaxation effects during the pulse may distort the shape of the excited volume.

2.5 Multiple-region MRI: Suppose only specific locations in the FOV contain signal, and the rest of the FOV is empty. Multiple-region MRI (mrMRI) allows these non-empty regions to be imaged without wasting time imaging all the empty spaces in between [11]. The smaller the non-empty re-

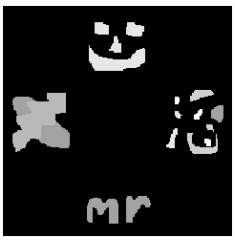

Fig. 4: Courtesy of Dr. Scott K. Nagle. Because it consists of 4 distinct parts, each one less than 1/9 of the FOV in size, this simulated object was reconstructed using only 4/9 of the k-space data points that would normally be required.

gions, and the fewer they are, the faster the acquisition can proceed. In the right circumstances, image acquisition can be accelerated by many-folds. Of course, the presence of signal in regions assumed to be empty would result in artifacts, and these artifacts would appear in the regions assumed to contain signal.

2.6 Reduced-FOV methods (Fig. 1b): In some applications, dynamic changes are expected mostly in a subset of the FOV. For example, in cardiac imaging, the beating heart is clearly much more dynamic than the tissues surrounding it. Reduced-FOV (rFOV) methods can accelerate acquisition by assuming that only a portion of the FOV is actually dynamic, the rest of the FOV being static [12,13]. If only one n^{th} of the FOV is dynamic, acquisition can be accelerated by up to a factor of n. Of course, artifacts are generated whenever motion occurs in a part of the FOV assumed to be static. This assumption can be softened using UNFOLD [14] (described in more details in 2.11), which allows temporally smooth changes to occur outside of the fully dynamic region. Noquist [15] is an elegant rFOV approach, which involves solving a large system of equations relating (unknown) pixel values at different time points to (measured) values in k-space. By assuming parts of the FOV to be static, the number of unknowns is reduced (static pixels have a same value at all time frames). The system becomes solvable if enough pixels can be assumed static [15].

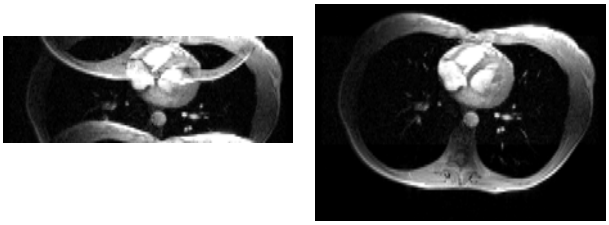

Fig. 5: The FOV is reduced by a factor of 2, reducing by 2 the number of required k-space lines. When reconstructed normally, images corrupted by aliasing are generated (left). Using UNFOLD and assuming the aliased material is not very dynamic, the de-aliased result shown on the right is obtained. The 7th frame of a 20-frame acquisition is shown here.

2.7 SVD and feature-recognizing MRI: Singular value decomposition (SVD) and feature-recognition are well-known topics in mathematics and computer science, with a range of applications that extends far beyond MRI. Using either SVD [16,17] or principal component analysis [18,19], one can extract a relatively small number of image components that strongly contribute in making, for example, a knee image look like a knee. By focusing on acquiring data relevant to only this small number of components, a knee image can be acquired extremely rapidly. In other words, one can use prior knowledge about how a given image should look like to achieve a reduction in data requirements. This prior knowledge may come either from analyzing a library of images acquired over many studies and many patients [18,19], or from one or more images acquired *in situ*, for a given patient [16,17]. The chosen "optimal" components can either be excited directly with specialized RF pulses [16,17], or decomposed into Fourier components and calculated from a subset of measured phase-encoding values [18,19]. Prior knowledge obtained through these methods can allow improvements in temporal resolution as a dynamic process (e.g., injection of contrast agent) is being imaged. Reductions in data requirements come at a cost: Features not present in the training image(s) may lead to artifacts, if they were to appear during the course of the dynamic acquisition. This limitation can be alleviated somewhat by modifying dynamically the acquisition strategy in light of the incoming data, as the dynamic acquisition is proceeding [20]. In summary, the lower the number of "optimal" components being used, the faster the acquisition can proceed, but the greater the risk of misrepresenting the object.

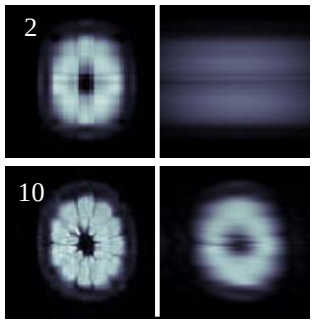

Fig. 6: Courtesy of Dr. W. Scott Hoge, Brigham and Women's Hospital, Harvard Medical School, Boston, USA. SVD encoding can generate images of an object using only a small amount of data. Using only 2 (top row) or 10 (bottom row) phase-encoded k-space lines led to dramatically blurred images (right). In contrast, using the same number of encodings, SVD imaging better captured the main features of the imaged object (left).

2.8 Keyhole (Fig. 1e): Keyhole imaging [21,22] assumes the following model:

object = (temporally dynamic, spatially smooth component) + (not very dynamic edge component)

The temporally dynamic, spatially smooth component is measured by acquiring repeatedly the k-space central region, with good temporal resolution (Fig. 1e). The edge information contained in the outer parts of k-space is acquired at one or two time points, before and/or after the dynamic acquisition (Fig. 1e) [21-24]. With one time point, the edge information is assumed static and is inserted in all of the (low spatial resolution) dynamic frames. With edge information both before and after the dynamic acquisition, temporal interpolation can be used to fill-in the missing data at each time frame. The smoother the dynamic component, the greater is the reduction in data requirements. Suppose the time changes are represented with a spatial resolution four times smaller than the full, reconstructed resolution. Only one fourth of the k-space lines would be required in the dynamic part of the acquisition, improving temporal resolution (i.e. the time between consecutive time frames) by a factor of 4. But as could be expected, dynamically changing edge information is not well captured, and may lead to image errors [23,24].

2.9 RIGR (Fig. 1e): RIGR [25,26] assumes a model of the form:

imaged object = (temporally dynamic, spatially smooth component) × (full spatial resolution component)

The equation above tends to oversimplify RIGR, but it should be sufficient to convey the main characteristics of the method. In RIGR, the data acquisition scheme is similar to that of keyhole: low-spatial resolution data is acquired with good temporal resolution, and full spatial-resolution data is acquired before and/or after the dynamic acquisition (Fig. 1e). Another similarity between keyhole and RIGR is the presence of a temporally dynamic, spatially smooth component in the model used to describe the imaged object. In fact, the reduction in data requirements achieved by these methods comes directly from the assumption that these terms are spatially smooth. Notice that the dynamic, spatially smooth term is additive in the keyhole model, but multiplicative in RIGR. This difference leads to different equations and algorithms to convert the limited acquired data set into a full one. RIGR assumes that the object at a given time point can be represented by multiplying a full resolution image with a smooth function. As for keyhole, the smoother this function, the greater is the reduction in data requirements. But of course, a smoother function also means a heavier assumption in the model, and thus a higher risk of artifacts.

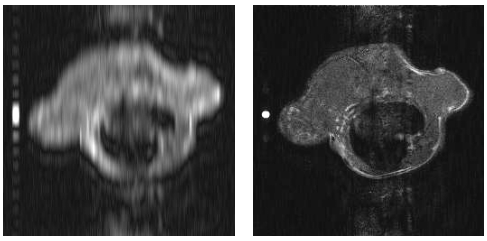

Fig. 7: Courtesy of Dr. Zhi-Pei Liang, University of Illinois at Urbana-Champaign, Urbana, USA. One time frame is displayed, from a dynamic study with injection of contrast agent in a rat. Only 16 k-space lines were acquired per time frame during the dynamic part of the acquisition. If the data were simply FFTed, a very blurred image would be obtained (left). RIGR can be used to estimate the edge information at this time frame, leading to a sharper image (right).

2.10 Sliding-window reconstruction (Fig. 1f): Fluoroscopy, or 'sliding-window reconstruction' [27], allows the number of reconstructed time frames to be increased by sharing data among neighboring 'true' time frames. For example, using the sampling scheme depicted in Fig. 1f, one could generate a full k-space matrix by combining the three first partially sampled matrices. A second time frame is generated by combining the partially sampled matrices 2, 3 and 4, a third time frame combines matrices 3, 4 and 5, etc. If each partially sampled matrix is acquired in a time τ, the reconstructed temporal resolution is τ (although the 'true' temporal resolution remains 3τ).

2.11 UNFOLD and k-t BLAST (Fig. 1f): Suppose one intentionally prescribes an FOV too small for the object, to speed-up the image acquisition process. Aliasing artifacts ensue, as several spatial locations may contribute to any given image pixel. If many time frames are acquired, the pixel value may change as a function of time, and the FFT of this time function gives a temporal frequency spectrum.

As shown in Ref. [14], applying a time-dependent k_y shift, $f(t)$, to the sampling function causes aliased signals to be modulated by a function $D(t) = \exp(i2\pi f(t)j)$, where j is the order of a given layer of aliasing

(j=0 for non-aliased signal), and where $f(t)$ is expressed as a fraction of the increment between k_y lines. In other words, one can control the phase of aliased signals by shifting the sampling function along k_y from time frame to time frame. Applications of UNFOLD published to date (including, for example, TSENSE [28], UNFOLD-SENSE [29,30], k-t BLAST [31,32], k-t SENSE [32]) all use the special case $f(t) = mt$, where t is the time-frame number and m is a proportionality constant, i.e., a slope. The appeal of this special case is that it makes $D(t)$ a Fourier function, and as a consequence, UNFOLD's effect in the Fourier domain becomes very simple: it displaces aliased signals, in the temporal frequency domain.

An example is depicted in Fig. 8a, where two components are present in the temporal frequency domain: a "true" non-aliased component, and one caused by aliasing. The ability to displace aliased signal, combined with assumption(s) about the width and/or shape of the corresponding spectra, may allow to efficiently stack them, somewhat like suitcases tightly packed in a car trunk (Fig. 8b). An UNFOLD fMRI example is depicted in Fig. 8c. In cardiac imaging, a reduction by almost a factor of 2 in data requirements can be obtained ([14], and reduced-FOV section above), and as much as a factor of 8 was obtained in fMRI [14]. UNFOLD can also be used to displace certain types of artifact signal around the temporal frequency domain, allowing such signal to be identified and removed [29].

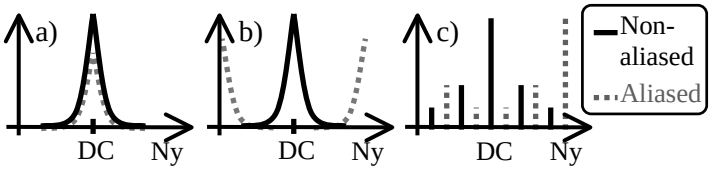

Fig. 8: Consider one pixel in a time series of images. The value of this pixel may change from time frame to time frame. Suppose one takes an FFT of this time function. a) The various aliased and non-aliased components overlapped in the image pixel also overlap in the temporal frequency domain. b) UNFOLD can displace aliased and non-aliased components with respect to each other. In this case, the aliased component was displaced all the way to the Nyquist frequency. c) Assuming fMRI data varies periodically in time with the period of the fMRI paradigm, a few spectra can be interleaved in the temporal frequency domain, allowing data to be packed efficiently (as much as 8 were interleaved in [14], although only 2 are depicted here).

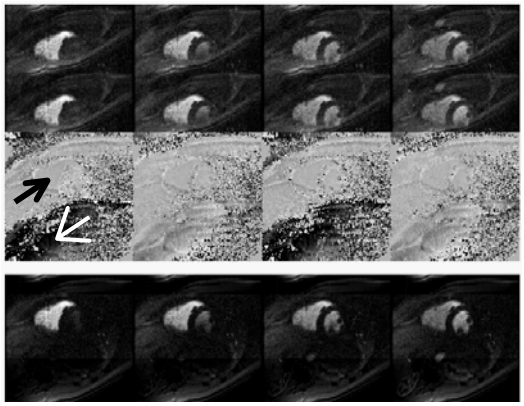

Fig. 9: Images corrupted by aliasing are acquired (top row, 4 consecutive time frames are shown). UNFOLD can force the aliased material to reverse its phase from frame to frame (middle row, white arrow), while non-aliased material does not (black arrow). Such Nyquist modulation of aliased signals was depicted in Fig. 8b. A simple time filter was used to suppress aliased signals (bottom row).

kt-BLAST [31,32] can be seen as an adaptive version of UNFOLD, as information from a training scan gets incorporated into the reconstruction. By using prior knowledge about the imaged object, the method allows imaging speed to be further increased. When methods such as UNFOLD or parallel imaging are pushed too far, they tend to generate ghosting artifacts. Imposing a resemblance with an artifact-free training dataset can be an effective way of suppressing such ghosting artifacts.

2.12 TRICKS: The TRICKS method [33,34] is tailored toward a 3D time-resolved angiography application. Like in keyhole, the k-space central region is sampled with good temporal resolution. But instead of sampling the outer parts of k-space only once or twice, TRICKS regularly samples these regions throughout the dynamic acquisition (although less often than it samples the central region). Linear temporal interpolation is used to evaluate all missing k-space locations, at all time frames. The reduction in data requirements comes from reducing the temporal resolution in the outer regions of k-space, and as a consequence, artifacts may arise if the edge content of the imaged object varies rapidly in time. Reductions by a few-fold in data requirements can be achieved, and 3D images can be generated with sufficient temporal resolution to resolve the arterial and venous phases, during dynamic enhancement.

3. HYBRID METHODS, OBTAINED THROUGH COMBINATIONS

The methods described above could be thought of as building blocks, to be assembled, mixed, and matched in a number of different ways. When different methods feed on different types of assumptions / prior knowledge, they can often be combined for added performance. For example, partial-Fourier imaging is the only one of these approaches to feed on Hermitian symmetries, and it has successfully been combined with most of the other methods discussed here (e.g., with parallel imaging [35], with UNFOLD [29]). In a somewhat similar fashion, because UNFOLD is a temporal strategy while many other methods function along spatial dimensions, it has proved compatible with several companion methods. More specifically, UNFOLD seems to have found a special niche as a helper to parallel imaging, leading to hybrid methods such as TSENSE [28] and UNFOLD-SENSE [29,30] (Fig. 10). As we have seen, the UNFOLD strategy can also be combined with a training-scan approach to give k-t BLAST [32], which in turn can be combined with parallel imaging to obtain k-t SENSE [32]. UNFOLD has also been combined with a 2D RF reduced-FOV approach, in a temperature-monitoring application [36]. Other examples of hybrid methods include a work by Mitsouras *et. al.* [37], where a training-scan approach was combined with parallel imaging, and a work by Parrish and Hu [38] where an rFOV approach was combined with keyhole, to allow spatially-smooth changes to occur outside of the dynamic region. Large acceleration factors have also been obtained combining TRICKS with a training-scan approach [39] (Fig. 11).

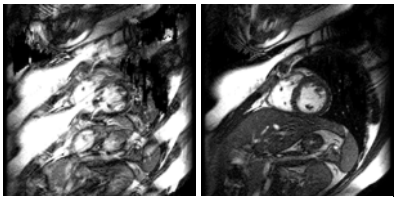

Fig. 10: When parallel imaging is pushed too far, it tends to generate ghosting artifacts (left, R = 3.0, 4-coil array). Adding UNFOLD can suppress these artifacts (right) [29].

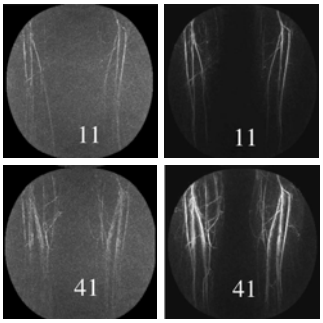

Fig. 11: Courtesy of Dr. Charles A. Mistretta, University of Wisconsin, Madison, USA. Two time frames (11th and 41st), in a time series of images, were reconstructed with PR TRICKS [34] (left). The HYPR strategy [39] injects prior knowledge about the imaged object into the reconstruction process. When added to PR TRICKS, HYPR can lead to significant improvements in image quality (right).

4. CONCLUSION

Once traditional means of increasing speed have been exhausted, the remaining option is to acquire only partial datasets, and filling-in the missing data using assumptions / prior knowledge. Many different methods, using a variety of different types of assumptions / prior knowledge, have been presented in the literature. The development, and also the combination, of partially sampled methods currently form a particularly active field of MR research. As a result, the available imaging speed has been rapidly increasing over the last few years, and can be expected to do so for some time, until fundamental SNR-related limits are reached, and impede further progress.

5. REFERENCES

1. Feinberg DA, Hale JD, Watts JC, Kaufman L, Mark A. Halving MR imaging time by conjugation: Demonstration at 3.5 kG. Radiology 1986;156:527-531.
2. Margosian P, Schmitt F, Purdy DE. Faster MR imaging: imaging with half the data. Healthcare Instr. 1986;1:195-197.
3. Noll DC, Nishimura DG, Macovski A. Homodyne detection in magnetic resonance imaging. IEEE Trans Med Imaging 1991;10:154-163.
4. Haacke EM, Lindskog ED, Lin W. A fast, iterative, partial-Fourier technique capable of local phase recovery. J Magn Reson 1991;92:126-145.
5. McGibney G, Smith MR, Nichols ST, Crawley A. Quantitative evaluation of several partial Fourier reconstruction algorithms used in MRI. Magn Reson Med 1993;30:51-59.
6. Peters DC, Korosec FR, Grist TM, Block WF, Holden JE, Vigen KK, Mistretta CA. Undersampled pro-

jection reconstruction applied to MR angiography. Magn Reson Med 2000;43:91-101.
7. Spielman DM, Pauly JM, Meyer CH. Magnetic resonance fluoroscopy using spirals with variable sampling densities. Magn Reson Med 1995;34:388-394.
8. Pipe JG. Reconstructing MR images from undersampled data: data-weighting considerations. Magn Reson Med 2000;43:867-875.
9. Pauly J, Nishimura D, Macovski A. A k-space analysis of small-tip-angle excitation. J Magn Reson 1989;81:43-56.
10. Pauly J, Le Roux P, Nishimura D, Macovski A. Parameter relations for the Shinnar-Le Roux selective excitation pulse design algorithm. IEEE Trans Med Imaging 1991;10:53-65.
11. Nagle SK, Levin DN. Multiple region MRI. Magn Reson Med 1999;41:774-786.
12. Hu X, Parrish T. Reduction of field of view for dynamic imaging. Magn Reson Med 1994;31:691-694.
13. Fredrickson JO, Pelc NJ. Temporal resolution improvement in dynamic imaging. Magn Reson Med 1996;35:621-625.
14. Madore B, Glover GH, Pelc NJ. Unaliasing by fourier-encoding the overlaps using the temporal dimension (UNFOLD), applied to cardiac imaging and fMRI. Magn Reson Med 1999;42:813-828.
15. Brummer ME, Moratal-Pérez D, Hong CY, Pettigrew RI, Millet-Roig J, Dixon WT. Noquist: reduced field-of-view imaging by direct Fourier inversion. Magn Reson Med 2004;51:331-342.
16. Zientara GP, Panych LP, Jolesz FA. Dynamically adaptive MRI with encoding by singular value decomposition. Magn Reson Med 1994;32:268-274.
17. Panych LP, Oesterle C, Zientara GP, Hennig J. Implementation of a fast gradient-echo SVD encoding technique for dynamic imaging. Magn Reson Med 1996;35:554-562.
18. Cao Y, Levin DN. Feature-recognizing MRI. Magn Reson Med 1993;30:305-317.
19. Cao Y, Levin DN. Using prior knowledge of human anatomy to constrain MR image acquisition and reconstruction: half k-space and full k-space techniques. Magn Reson Imaging 1997;15:669-677.
20. Hoge WS, Miller EL, Lev-Ari H, Brooks DH, Panych LP. A doubly adaptive approach to dynamic MRI sequence estimation. IEEE TIP 2002;11:1168-1178.
21. van Vaals JJ, Brummer ME, Dixon WT, Tuithof HH, Engels H, Nelson RC, Gerety BM, Chezmar JL, den Boer JA. "Keyhole" method for accelerating imaging of contrast agent uptake. J Magn Reson Imaging 1993;3:671-675.
22. Jones RA, Haraldseth O, Müller TB, Rinck PA, Øksendal AN. K-space substitution: a novel dynamic imaging technique. Magn Reson Med 1993;29:830-834.
23. Bishop JE, Santyr GE, Kelcz F, Plewes DB. Limitations of the keyhole technique for quantitative dynamic contrast-enhanced breast MRI. J Magn Reson Imaging 1997;7:716-723.
24. Oesterle C, Strohschein R, Köhler M, Schnell M, Hennig J. Benefits and pitfalls of keyhole imaging, especially in first-pass perfusion studies. J Magn Reson Imaging 2000;11:312-323.
25. Liang ZP, Lauterbur PC. An efficient method for dynamic magnetic resonance imaging. IEEE Trans Med Imaging 1994;13:677-686.
26. Hanson JM, Liang ZP, Wiener EC, Lauterbur PC. Fast dynamic imaging using two reference images. Magn Reson Med 1996;36:172-175.
27. Riederer SJ, Tasciyan T, Farzaneh F, Lee JN, Wright RC, Herfkens RJ. MR fluoroscopy: technical feasibility. Magn Reson Med 1988;8:1-15.
28. Kellman P, Epstein FH, McVeigh ER. Adaptive sensitivity encoding incorporating temporal filtering (TSENSE). Magn Reson Med 2001;45:846-852.
29. Madore B. Using UNFOLD to remove artifacts in parallel imaging and in partial-Fourier imaging. Magn Reson Med 2002;48:493-501.
30. Madore B. UNFOLD-SENSE: a parallel MRI method with self-calibration and artifact suppression. Magn Reson Med 2004;52:310-320.
31. Tsao J. On the UNFOLD method. Magn Reson Med 2002;47:202-207.
32. Tsao J, Boesiger P, Pruessmann KP. k-t BLAST and k-t SENSE: dynamic MRI with high frame rate exploiting spatiotemporal correlations. Magn Reson Med 2003;50:1031-1042.
33. Korosec FR, Frayne R, Grist TM, Mistretta CA. Time-resolved contrast-enhanced 3D MR angiography. Magn Reson Med 1996;36:345-351.
34. Vigen KK, Peters DC, Grist TM, Block WF, Mistretta CA. Undersampled projection-reconstruction imaging for time-resolved contrast-enhanced imaging. Magn Reson Med 2000;43:170-176.
35. Bydder M, Robson MD. Partial fourier partially parallel imaging. Magn Reson Med 2005;53:1393-1401.
36. Zhao L, Madore B, Panych LP. Reduced field-of-view MRI with two-dimensional spatially-selective RF excitation and UNFOLD. Magn Reson Med 2005;53:1118-1125.
37. Mitsouras D, Hoge WS, Rybicki FJ, Kyriakos WE, Edelman A, Zientara GP. Non-Fourier-encoded parallel MRI using multiple receiver coils. Magn Reson Med 2004;52:321-328.
38. Parrish TB, Hu X. Hybrid technique for dynamic imaging. Magn Reson Med 2000;44:51-55.
39. Mistretta CA, Wieben O, Velikina J, Block W, Perry J, Wu Y, Johnson K, Wu Y. Highly constrained backprojection for time-resolved MRI. Accepted for publication in Magn Reson Med.